arXiv:1702.05859v1 [math.NA] 20 Feb 2017

DATA-DRIVEN POLYNOMIAL RIDGE APPROXIMATION USING VARIABLE PROJECTION*

JEFFREY M. HOKANSON† AND PAUL G. CONSTANTINE†

Abstract. Inexpensive surrogates are useful for reducing the cost of science and engineering studies with large-scale computational models that contain many input parameters. A ridge approximation is a surrogate distinguished by its model form: namely, a nonlinear function of a few linear combinations of the input parameters. Parameter studies (e.g., optimization or uncertainty quantification) with ridge approximations can exploit its low-dimensional structure by working on the coordinates of the subspace defined by the linear combination weights, reducing the effective dimension. We introduce a new, fast algorithm for constructing a least-squares-fit polynomial ridge approximation from function samples. Naively, this would require optimizing both the polynomial coefficients and the subspace. However, given a fixed subspace the optimal polynomial coefficients solve a linear least-squares problem. Our proposed method exploits this structure by implicitly computing these coefficients using *variable projection*, which leaves an optimization problem over the subspace alone. We present an algorithm that finds the optimal subspace by optimizing over the Grassmann manifold using a Gauss-Newton Hessian approximation. We provide the details of the optimization algorithm, and we demonstrate its performance on several numerical examples. The Gauss-Newton method has superior theoretical guarantees and faster convergence than an alternating heuristic for ridge approximation proposed by Constantine, Eftekhari, and Ward [*arXiv* 1606.01929] that (i) optimizes the polynomial coefficients given the subspace and (ii) optimizes the subspace given the coefficients.

Key words. active subspaces, ridge function, variable projection

AMS subject classifications. 49M15, 62J02, 90C53

1. Introduction.

We consider a scalar-valued function of several variables denoted $f(\mathbf{x})$ with $\mathbf{x} \in \mathcal{D} \subseteq \mathbb{R}^m$. The functions of interest come from large-scale computational science and engineering models, where the variables $\mathbf{x}$ represent physical input parameters and f represents an output quantity of interest. Such functions are typically expensive to evaluate for a given $\mathbf{x}$ since they involve numerically approximating the solution of governing equations that model the physical system. Given $f(\mathbf{x})$, the engineer may seek to optimize f over $\mathbf{x}\mathcal{D}$ or characterize the uncertainty in f given a model for uncertainty in $\mathbf{x}$. Both of these problems require evaluating $f(\mathbf{x})$ many times. To reduce the cost of such studies, we seek to replace this expensive function with an inexpensive surrogate—a practice that is common in both optimization [23, 33] and uncertainty quantification [34, Ch. 13]. The challenge is that, in the absence of known exploitable structure, the cost of constructing an accurate approximation scales exponentially with the number input parameters—a phenomenon known as the *curse of dimensionality* [28]. Even ignoring accuracy, a polynomial surrogate of degree p on the m dimensional space requires at least $\mathcal{O}(p^m)$ samples of $f(\mathbf{x})$ for its construction to be well posed as $m \to \infty$. To avoid this problem, we must assume f has some exploitable low-dimensional structure. Recent experience suggests that many engineering problems have a low-dimensional *active subspace* [6], which is the span of the set of important directions in the space of the input parameters; perturbing the input parameters along these directions changes f more, on average, than perturbing the input parameters orthogonally to the active

*Submitted to the editors 20 February 2017.

Funding: This work was supported by the Department of Defense, Defense Advanced Research Project Agency's program Enabling Quantification of Uncertainty in Physical Systems.

† Department of Applied Mathematics and Statistics, Colorado School of Mines, 1500 Illinois St, Golden CO, 80401, ({hokanson, pconstan}@mines.edu).

subspace. Such structure has been identified in integrated hydrologic models [22], a solar cell circuit model [11], a subsurface permeability model [18], a lithium-ion battery model [7], a magnetohydrodynamics power generation model [19], a hypersonic scramjet model [10], an annular combustor model [3], aerospace vehicles [25, 26], and automobiles [29]. Consequently, our goal is to construct a surrogate of $f(\mathbf{x})$ based on the ridge approximation:

$$f(\mathbf{x}) \approx g(\mathbf{U}^\top \mathbf{x}), \quad \text{where} \quad \mathbf{U} \in \mathbb{R}^{m\times n} \text{ and } \mathbf{U}^\top\mathbf{U} = \mathbf{I}. \tag{1}$$

Here we construct this approximation in a point-wise sense, seeking to minimize the ℓ_2 error of approximating $f(\mathbf{x})$ by $g(\mathbf{U}^\top\mathbf{x})$ for a given set of points $\{\mathbf{x}_i\}_{i=1}^N \subset \mathcal{D}$

$$\underset{\substack{g\in\mathbb{P}^p(\mathbb{R}^n)\\ \text{Range}\,\mathbf{U}\in\mathbb{G}(n,\mathbb{R}^m)}}{\text{minimize}} \quad \sum_{i=1}^N \left[f(\mathbf{x}_i) - g(\mathbf{U}^\top\mathbf{x}_i)\right]^2; \tag{2}$$

where $\mathbb{P}^p(\mathbb{R}^n)$ denotes the space of polynomials on $\mathbb{R}^n$ of total degree p and $\mathbb{G}(n,\mathbb{R}^m)$ is the Grassmann manifold of n-dimensional subspaces of $\mathbb{R}^m$ which we non-uniquely parameterize by matrices $\mathbf{U} \in \mathbb{R}^{m\times n}$ with orthonormal columns following [15, §2.1] Hence, this *polynomial ridge approximation* requires only $\mathcal{O}(mn + p^n)$ samples for construction—exponentially fewer than the $\mathcal{O}(p^m)$ samples required for the equivalent polynomial model on all m input parameters. We call the surrogate constructed by solving (2) a *data-driven polynomial ridge approximation* since it is constructed using only samples $f(\mathbf{x}_i)$ which, for example, may come from a complex mathematical model or a physical experiment. In this paper we develop an efficient algorithm for solving the data-driven polynomial ridge approximation (2) using *variable projection* [20] to implicitly solve for the polynomial g and solve the remaining optimization problem over the subspace spanned by $\mathbf{U}$ using a Gauss-Newton method developed following Edelman, Arias, and Smith [15].

In the vast literature on surrogate models, and among ridge functions in particular, our polynomial ridge approximation is comparatively rare. More common is the model used by *projection pursuit regression* [13, 17]:

$$f(\mathbf{x}) \approx \alpha + \sum_{j=1}^n g_j(\mathbf{a}_j^\top\mathbf{x}), \qquad \mathbf{a}_j \in \mathbb{R}^m, \tag{3}$$

where g_j is a *smoothing function*, such a logistic function or a spline [24, 31]. If g_j is taken to be a polynomial, this resembles our construction but misses the cross terms inherent in a polynomial over n-dimensions. Some recent work focuses on finding the matrix $\mathbf{U} \in \mathbb{R}^{m\times n}$ whose range defines the ridge subspace using compressive sensing and then separately finds g by polynomial interpolation; see, e.g., [16, 36]. However, these works require the ability to evaluate f at arbitrary points to estimate finite difference derivatives. The only work that explicitly considers our problem with a fixed sample of points is Algorithm 2 of Constantine, Eftekhari, and Ward [9], which solves (2) using an alternating iteration—switching between optimizing the subspace defined by $\mathbf{U}$ and the polynomial approximation g. In this paper we improve on their work, presenting an algorithm for the data-driven polynomial ridge approximation that is up to four thousand times faster and significantly more accurate in our numerical examples. This speedup is in part due to a better implementation that exploits the structure of the problem and in part due to the quadratic convergence of Gauss-Newton for zero-residual problems as opposed to the linear convergence of alternating methods [32, §3].

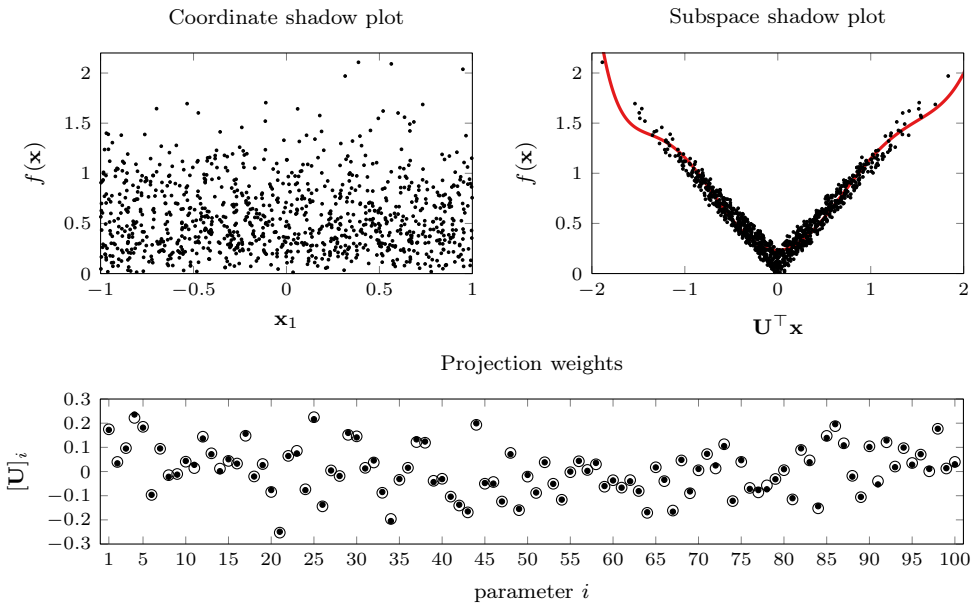

FIG. 1. *Ridge approximations reveal structure not present in coordinate perspectives. In this example,* $f : \mathcal{D} := [-1,1]^{100} \to \mathbb{R}$ *with* $f(\mathbf{x}) = |\widehat{\mathbf{u}}^\top\mathbf{x}| + 0.1(\sin(1000[\mathbf{x}]_2) + 1)$ *where* $\widehat{\mathbf{u}}$ *has been sampled uniformly on the unit sphere and the sine term simulates deterministic noise. Viewed from a coordinate perspective on the upper left, no structure is present. However, after fitting a ridge function to* $N = 1000$ *samples of* f *with polynomial degree* $p = 7$ *and subspace dimension* $n = 1$ *using our Algorithm* 1*, the structure of* f *is revealed by looking along the recovered subspace* $\mathbf{U}$*. On the bottom, we see that the coefficients of* $\mathbf{U}$*,* $\bullet$*, closely match the coefficients of* $\widehat{\mathbf{u}}$*,* $\circ$*.*

In the remainder of this paper we first reformulate the polynomial ridge approximation as a separable nonlinear least squares problem in section 2. This permits us to remove the polynomial coefficients and solve over the subspace defined by $\mathbf{U}$ alone using variable projection [20]. Then in section 3 we develop a Gauss-Newton algorithm for optimizing the subspace spanned by $\mathbf{U}$ on the Grassmann manifold. Section 4 then provides examples comparing our algorithm to the alternating approach of Constantine, Eftekhari, and Ward. In addition to toy problems with a polynomial f, we also demonstrate the effectiveness of our approach by constructing polynomial ridge approximations of two f from application problems: an 18-dimensional airfoil model [6, §5.3.1] and a 100-dimensional elliptic PDE problem [8]. Finally, we conclude with a discussion on selecting the polynomial order and subspace dimension, how an iterative scheme might be devised to select the optimum sample points $\mathbf{x}_i$, and how the numerical stability might be improved.

2. Separable reformulation. The key to efficiently solving the data-driven polynomial ridge approximation problem is reducing the number parameters we optimize over. This approximation problem in (2) requires minimizing over both the polynomial g and the subspace spanned by $\mathbf{U}$. In this section, we reformulate (2) exposing its structure as a *separable nonlinear least squares problem*. This permits us to implicitly solve for the polynomial g, leaving an optimization over the subspace spanned by $\mathbf{U}$ alone. We obtain the Jacobian of this smaller problem using *variable projection* [20] which enables us to efficiently solve the remaining optimization as outlined in section 3.

To observe the separable structure inherent in the ridge approximation problem, first note that any polynomial in $\mathbb{P}^p(\mathbb{R}^n)$ can be written in the monomial basis as a sum over the multi-index $\boldsymbol{\alpha} \in \mathbb{N}^n$,

$$g(\mathbf{y}) = \sum_{|\boldsymbol{\alpha}| \leq p} c_{\boldsymbol{\alpha}} \mathbf{y}^{\boldsymbol{\alpha}} = \sum_{j=1}^{\binom{n+p}{p}} c_{[\boldsymbol{\alpha}]_j} \prod_{k=1}^{n} y_k^{[\boldsymbol{\alpha}_j]_k}. \tag{4}$$

Using this construction, we rewrite $g(\mathbf{U}^\top \mathbf{x})$ as the matrix vector product:

$$g(\mathbf{U}^\top \mathbf{x}_i) = [\mathbf{V}(\mathbf{U})\mathbf{c}]_i, \qquad [\mathbf{V}(\mathbf{U})]_{i,j} := \prod_{k=1}^{n} [\mathbf{U}^\top \mathbf{x}_i]^{[\boldsymbol{\alpha}_j]_k}, \qquad [\mathbf{c}]_j = c_{[\alpha]_j} \tag{5}$$

where $\mathbf{V}(\mathbf{U})$ is a multivariate Vandermonde matrix. This yields a more compact representation of the minimization problem that exposes the separable structure:

$$\underset{\substack{g \in \mathbb{P}^p(\mathbb{R}^n) \\ \text{Range}\, \mathbf{U} \in \mathbb{G}(n, \mathbb{R}^m)}}{\text{minimize}} \sum_{i=1}^{N} \left[f(\mathbf{x}_i) - g(\mathbf{U}^\top \mathbf{x}_i) \right]^2 \quad \Longleftrightarrow \quad \underset{\substack{\mathbf{c} \in \mathbb{R}^{\binom{n+p}{p}} \\ \text{Range}\, \mathbf{U} \in \mathbb{G}(n, \mathbb{R}^m)}}{\text{minimize}} \quad \|\mathbf{f} - \mathbf{V}(\mathbf{U})\mathbf{c}\|_2^2, \tag{6}$$

where $[\mathbf{f}]_i = f(\mathbf{x}_i)$. This second minimization problem is called *separable* as the optimization parameters have been partitioned into two classes: the linear parameters $\mathbf{c}$ that are easy to solve for given $\mathbf{U}$ and the nonlinear parameters $\mathbf{U}$ that are more challenging to solve for. In the next subsection we use variable projection to implicitly remove the linear parameters $\mathbf{c}$ and conclude with an orthogonality result concerning the Jacobian and $\mathbf{U}$ in subsection 2.2.

2.1. Variable projection. The key insight of variable projection is that if the nonlinear parameters $\mathbf{U}$ are fixed, the linear parameters $\mathbf{c}$ are easily recovered using the pseudoinverse:

$$\underset{\mathbf{c}}{\text{minimize}} \ \|\mathbf{f} - \mathbf{V}(\mathbf{U})\mathbf{c}\|_2 \quad \Rightarrow \quad \mathbf{c} = \mathbf{V}(\mathbf{U})^+ \mathbf{f}. \tag{7}$$

Then, replacing $\mathbf{c}$ with $\mathbf{V}(\mathbf{U})^+\mathbf{f}$ in (6) reveals an optimization problem over $\mathbf{U}$ alone:

$$\underset{\text{Range}\, \mathbf{U} \in \mathbb{G}(n, \mathbb{R}^m)}{\text{minimize}} \ \|\mathbf{f} - \mathbf{V}(\mathbf{U})\mathbf{V}(\mathbf{U})^+ \mathbf{f}\|_2^2. \tag{8}$$

Recognizing $\mathbf{V}(\mathbf{U})\mathbf{V}(\mathbf{U})^+$ as an orthogonal projector onto the range of $\mathbf{V}(\mathbf{U})$, we can rewrite the minimization problem as

$$\underset{\text{Range}\, \mathbf{U} \in \mathbb{G}(n, \mathbb{R}^m)}{\text{minimize}} \ \|\mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})} \mathbf{f}\|_2^2 \tag{9}$$

where $\mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})}$ is the orthogonal projector onto the complement of the range of $\mathbf{V}(\mathbf{U})$.

Golub and Pereyra reveal a formula for the derivative of the residual $\mathbf{r}(\mathbf{U}) = \mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})}\mathbf{f}$ [20]. Denoting the derivative of the residual as a tensor $\boldsymbol{\mathcal{J}}(\mathbf{U}) \in \mathbb{R}^{N \times m \times n}$ where $[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,j,k} = \partial [\mathbf{r}(\mathbf{U})]_i / \partial [\mathbf{U}]_{j,k}$, then

$$[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{\cdot,j,k} = -\left[\left(\mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})} \frac{\partial \mathbf{V}}{\partial [\mathbf{U}]_{j,k}}(\mathbf{U}) \ \mathbf{V}(\mathbf{U})^- \right) + \left(\mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})} \frac{\partial \mathbf{V}}{\partial [\mathbf{U}]_{j,k}}(\mathbf{U}) \ \mathbf{V}(\mathbf{U})^- \right)^\top \right] \mathbf{f}, \tag{10}$$

where $^-$ denotes a least squares pseudoinverse such that $\mathbf{V}\mathbf{V}^-\mathbf{V} = \mathbf{V}$ and $\mathbf{V}\mathbf{V}^- = (\mathbf{V}\mathbf{V}^-)^\top$. Then, from the definition of the multivariate Vandermonde matrix $\mathbf{V}(\mathbf{U})$ in (5), its derivative is:

$$\left[\frac{\partial \mathbf{V}}{\partial [\mathbf{U}]_{k,\ell}}(\mathbf{U})\right]_{i,j} = [\boldsymbol{\alpha}_j]_\ell \, [\mathbf{U}^\top \mathbf{x}_i]_\ell^{[\boldsymbol{\alpha}_j]_\ell - 1} \, [\mathbf{x}_i]_k \prod_{\substack{q=1 \\ q\neq\ell}}^{n} [\mathbf{U}^\top \mathbf{x}_i]_q^{[\boldsymbol{\alpha}_j]_q} \, . \tag{11}$$

2.2. Orthogonality. In addition to constructing explicit formula for the Jacobian of $\mathbf{r}(\mathbf{U})$, we also show that slices of this Jacobian $\boldsymbol{\mathcal{J}}(\mathbf{U})$ are orthogonal to $\mathbf{U}$, $\mathbf{U}^\top[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot} = \mathbf{0}$. This property reduces the computational burden during optimization as later discussed in subsection 3.2.

THEOREM 1. *If $\boldsymbol{\mathcal{J}}(\mathbf{U})$ is defined as in* (10)*, then $\mathbf{U}^\top[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot} = \mathbf{0}$ $\forall i$.*

Proof. First note,

$$\begin{aligned}\left[\mathbf{U}^\top[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}\right]_{j,k} &= \sum_\ell [\mathbf{U}]_{\ell,j}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\ell,k} \\ &= -\sum_\ell [\mathbf{U}]_{\ell,j}\left[\mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\frac{\partial \mathbf{V}(\mathbf{U})}{\partial [\mathbf{U}]_{\ell,k}}\mathbf{V}(\mathbf{U})^{-1}\mathbf{f} + \mathbf{V}(\mathbf{U})^{-\top}\frac{\partial \mathbf{V}(\mathbf{U})}{\partial [\mathbf{U}]_{\ell,k}}^\top \mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\mathbf{f}\right]_i .\end{aligned}$$

By linearity, we take the sum inside each term, leaving

$$= -\left[\mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\left[\sum_\ell [\mathbf{U}]_{\ell,j}\frac{\partial \mathbf{V}(\mathbf{U})}{\partial [\mathbf{U}]_{\ell,k}}\right]\mathbf{V}(\mathbf{U})^{-1}\mathbf{f} + \mathbf{V}(\mathbf{U})^{-\top}\left[\sum_\ell [\mathbf{U}]_{\ell,j}\frac{\partial \mathbf{V}(\mathbf{U})}{\partial [\mathbf{U}]_{\ell,k}}^\top\right]\mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\mathbf{f}\right]_i . \tag{12}$$

Defining this interior sum as the matrix $\mathbf{W}$, we note $\mathbf{W}$ is in the range of $\mathbf{V}(\mathbf{U})$:

$$\begin{aligned}\mathbf{W} := \left[\sum_\ell [\mathbf{U}]_{\ell,j}\frac{\partial \mathbf{V}(\mathbf{U})}{\partial [\mathbf{U}]_{\ell,k}}\right]_{r,q} &= \sum_\ell [\mathbf{U}]_{\ell,j}[\boldsymbol{\alpha}_q]_k [\mathbf{U}^\top \mathbf{x}_r]_k^{[\boldsymbol{\alpha}_q]_k - 1}[\mathbf{x}_r]_\ell \prod_{\substack{s=1 \\ s\neq k}}^{n} [\mathbf{U}^\top \mathbf{x}_r]_s^{[\boldsymbol{\alpha}_q]_s} \\ &= [\boldsymbol{\alpha}_q]_k [\mathbf{U}^\top \mathbf{x}_r]_\ell^{[\boldsymbol{\alpha}_q]_k - 1}[\mathbf{U}^\top \mathbf{x}_r]_j \prod_{\substack{s=1 \\ s\neq k}}^{n} [\mathbf{U}^\top \mathbf{x}_r]_s^{[\boldsymbol{\alpha}_q]_s}.\end{aligned}$$

Thus, as $\mathbf{W} \in \text{Range}(\mathbf{V}(\mathbf{U}))$, then $\mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\mathbf{W} = \mathbf{0}$. Then as this product appears in both terms of (12), we conclude $\mathbf{U}^\top[\boldsymbol{\mathcal{J}}]_{i,\cdot,\cdot} = \mathbf{0}$. ∎

3. Optimization on the Grassmann manifold. Having implicitly found the polynomial g using variable projection in the previous section, we now develop an algorithm for solving the optimization problem over $\mathbf{U}$ alone

$$\underset{\text{Range}\,\mathbf{U}\in\mathbb{G}(n,\mathbb{R}^m)}{\text{minimize}} \; \|\mathbf{P}^\perp_{\mathbf{V}(\mathbf{U})}\mathbf{f}\|_2^2. \tag{13}$$

Optimization over the Grassmann manifold of n-dimensional subspaces of $\mathbb{R}^m$ is more complex than optimization on Euclidean space. Here we follow the approach of Edelman, Arias, and Smith [15] where $\mathbf{U}$ satisfies the implicit constraint $\mathbf{U}^\top\mathbf{U} = \mathbf{I}$. We first review Newton's method on the Grassmann manifold following their construction before modifying their approach to construct a Gauss-Newton method on the Grassmann manifold for the data-driven polynomial ridge approximation problem.

3.1. Newton's method on the Grassmann manifold. To begin, we first review Newton's method on the Grassmann manifold, following Edelman, Arias, and Smith [15]. There are two key properties we consider: how to update $\mathbf{U}$ given a search direction $\mathbf{\Delta}$ and how to choose the search direction $\mathbf{\Delta}$ using Newton's method.

When optimizing in a Euclidean space, given a direction $\mathbf{d}$ and an initial point $\mathbf{x}_0$, the next iterate is chosen along the trajectory $\mathbf{x}(t) = \mathbf{x}_0 + t\mathbf{d}$ for $t \in (0, \infty)$ where t is selected to ensure convergence. However, if we were to apply the same search strategy on our parameterization Grassmann manifold, this would result in a point that does not obey the orthogonality constraint: $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$. Instead, we replace the linear trajectory with by a geodesic (a contour with constant derivative). Then, following [15, Thm. 2.3], if $\mathbf{\Delta} \in \mathbb{R}^{m\times n}$ is the search direction that is tangent to the current estimate $\mathbf{U}_0$, $\mathbf{U}_0^\top \mathbf{\Delta} = \mathbf{0}$, then we choose $\mathbf{U}(t)$ on the geodesic

$$\mathbf{U}(t) = \mathbf{U}_0 \mathbf{Z} \cos(\mathbf{\Sigma} t)\mathbf{Z}^\top + \mathbf{Y} \sin(\mathbf{\Sigma} t)\mathbf{Z}^\top \tag{14}$$

where $\mathbf{\Delta} = \mathbf{Y}\mathbf{\Sigma}\mathbf{Z}^\top$ is the short form singular value decomposition.

In addition to changing the search trajectories, optimization on the Grassmann manifold changes the gradient and Hessian. Here we consider an arbitrary function ϕ of a matrix $\mathbf{U} \in \mathbb{R}^{m\times n}$ with orthonormal columns. Then we define the first and second derivatives of ϕ with respect to the entries of $\mathbf{U}$ following [15, (2.52)&(2.56)]:

$$[\phi_{\mathbf{U}}]_{i,j} := \frac{\partial \phi}{\partial [\mathbf{U}]_{i,j}}, \qquad \phi_{\mathbf{U}} \in \mathbb{R}^{m\times n}; \tag{15}$$

$$[\phi_{\mathbf{U}\mathbf{U}}]_{i,j,k,\ell} := \frac{\partial^2 \phi}{\partial [\mathbf{U}]_{i,j} \partial [\mathbf{U}]_{k,\ell}}, \qquad \phi_{\mathbf{U}} \in \mathbb{R}^{m\times n\times m\times n}. \tag{16}$$

Using these definitions, the gradient of ϕ on the Grassmann manifold is [15, (2.70)]

$$\operatorname{grad} \phi = \phi_{\mathbf{U}} - \mathbf{U}\mathbf{U}^\top \phi_{\mathbf{U}} = \mathbf{P}_{\mathbf{U}}^{\perp} \phi_{\mathbf{U}} \in \mathbb{R}^{m\times n}. \tag{17}$$

The Hessian of ϕ on the Grassmann manifold, $\operatorname{Hess} \phi \in \mathbb{R}^{m\times n\times m\times n}$, is defined by its action on two test matrices $\mathbf{\Delta}, \mathbf{X} \in \mathbb{R}^{m\times n}$ [15, (2.71)]

$$\operatorname{Hess} \phi(\mathbf{\Delta}, \mathbf{X}) = \sum_{i,j,k,\ell} [\phi_{\mathbf{U}\mathbf{U}}]_{i,j,k,\ell} [\mathbf{\Delta}]_{i,j} [\mathbf{X}]_{k,\ell} - \operatorname{Tr}(\mathbf{\Delta}^\top \mathbf{X} \mathbf{U}^\top \phi_{\mathbf{U}}). \tag{18}$$

Using these definitions, Newton's method on the Grassmann manifold at $\mathbf{U}$ chooses the tangent search direction $\mathbf{\Delta} \in \mathbb{R}^{m\times n}$ satisfying [15, (2.58)]

$$\mathbf{U}^\top \mathbf{\Delta} = \mathbf{0} \quad \text{and} \quad \operatorname{Hess} \phi(\mathbf{\Delta}, \mathbf{X}) = -\langle \operatorname{grad} \phi, \mathbf{X} \rangle \quad \forall\ \mathbf{X} \text{ such that } \mathbf{U}^\top \mathbf{X} = 0 \tag{19}$$

where the inner product $\langle \cdot, \cdot \rangle$ on this space is [15, (2.69)]

$$\langle \mathbf{X}, \mathbf{Y} \rangle = \operatorname{Tr} \mathbf{X}^\top \mathbf{Y} = (\operatorname{vec} \mathbf{X})^\top (\operatorname{vec} \mathbf{Y}) \tag{20}$$

and vec maps the matrix $\mathbb{R}^{m\times n}$ into the vector $\mathbb{R}^{mn}$.

3.2. Grassmann Gauss-Newton. The challenge with Newton's method is it requires second derivative information which is hard to obtain for our problem. Here we replace the true Hessian with the Gauss-Newton approximation, yielding a Grassmann Gauss-Newton method. We are not the first to construct a Grassmann Gauss-Newton method, c.f., [1, §8.4]; however, we are able to exploit the structure of the polynomial ridge approximation problem to simplify computation. Via Theorem 1 we

can replace the normal equations (19) with a better conditioned least squares problem (32), analogously to the Gauss-Newton method in Euclidean space [27, (10.26)].

To begin, our data-driven polynomial ridge approximation objective function is

$$\phi(\mathbf{U}) = \frac{1}{2}\|\mathbf{P}^{\perp}_{\mathbf{V}(\mathbf{U})}\mathbf{f}\|_2^2 = \frac{1}{2}\|\mathbf{r}(\mathbf{U})\|_2^2. \tag{21}$$

Then the first and second derivatives of ϕ are

$$\phi_{\mathbf{U}} = \sum_i [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}[\mathbf{r}(\mathbf{U})]_i; \tag{22}$$

$$[\phi_{\mathbf{UU}}]_{i,j,k,\ell} = \sum_q [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,i,j}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,k,\ell} + \sum_q [\mathbf{r}(\mathbf{U})]_q \frac{[\mathbf{r}(\mathbf{U})]_q}{\partial[\mathbf{U}]_{i,j}\partial[\mathbf{U}]_{k,\ell}}. \tag{23}$$

Invoking the Gauss-Newton approximation, we drop the second term above

$$[\phi_{\mathbf{UU}}]_{i,j,k,\ell} \approx [\widetilde{\phi}_{\mathbf{UU}}]_{i,j,k,\ell} := \sum_q [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,i,j}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,k,\ell}. \tag{24}$$

Then replacing $\phi_{\mathbf{UU}}$ with $\widetilde{\phi}_{\mathbf{UU}}$ in the Hessian (18) yields the approximate Hessian

$$\widetilde{\text{Hess}}\ \phi(\boldsymbol{\Delta}, \mathbf{X}) := \sum_{i,j,k,\ell,q} [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,i,j}[\boldsymbol{\Delta}]_{i,j}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,k,\ell}[\mathbf{X}]_{k,\ell} - \text{Tr}(\boldsymbol{\Delta}^\top \mathbf{X}\mathbf{U}^\top \phi_{\mathbf{U}}). \tag{25}$$

Immediately, we note that by Theorem 1 $\mathbf{U}^\top\phi_{\mathbf{U}}$ is zero:

$$\mathbf{U}^\top\phi_{\mathbf{U}} = \sum_i \mathbf{U}^\top[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}[\mathbf{r}(\mathbf{U})]_i = \mathbf{0}. \tag{26}$$

Hence, the second term drops out of the approximate Hessian, leaving

$$\widetilde{\text{Hess}}\ \phi(\boldsymbol{\Delta}, \mathbf{X}) = \sum_{i,j,k,\ell,q} [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,i,j}[\boldsymbol{\Delta}]_{i,j}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{q,k,\ell}[\mathbf{X}]_{k,\ell}. \tag{27}$$

This summation can be rearranged to look like the more familiar Hessian approximation $\mathbf{J}^\top\mathbf{J}$ when the Jacobian $\mathbf{J}$ is a matrix. Defining the vectorization operator for tensors that maps $\boldsymbol{\mathcal{J}}(\mathbf{U}) \in \mathbb{R}^{N\times m\times n}$ to a matrix in $\mathbb{R}^{N\times mn}$, this product above is

$$\widetilde{\text{Hess}}\ \phi(\boldsymbol{\Delta}, \mathbf{X}) = (\text{vec}\,\mathbf{X})^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))(\text{vec}\,\boldsymbol{\Delta}). \tag{28}$$

With this familiar expression for the approximate Hessian, we now seek to rework the Newton step from a square linear system into an overdetermined least squares problem. First, by invoking Theorem 1 through (26), we note

$$\text{grad}\,\phi(\mathbf{U}) = \phi_{\mathbf{U}} - \mathbf{U}\mathbf{U}^\top\phi_{\mathbf{U}} = \phi_{\mathbf{U}} = (\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))^\top\mathbf{r}(\mathbf{U}). \tag{29}$$

Then, examining the right hand side of the Newton step (19),

$$\langle \text{grad}\,\phi, \mathbf{X}\rangle = (\text{vec}\,\mathbf{X})^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))^\top\mathbf{r}(\mathbf{U}).$$

Thus, the Gauss-Newton step $\boldsymbol{\Delta} \in \mathbb{R}^{m\times n}$ is the matrix satisfying $\mathbf{U}^\top\boldsymbol{\Delta} = \mathbf{0}$ and

$$(\text{vec}\,\mathbf{X})^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))(\text{vec}\,\boldsymbol{\Delta}) = (\text{vec}\,\mathbf{X})^\top(\text{vec}\,\boldsymbol{\mathcal{J}}(\mathbf{U}))^\top\mathbf{r}(\mathbf{U}), \forall \mathbf{X} \text{ s.t. } \mathbf{U}^\top\mathbf{X} = \mathbf{0}.$$

To convert this into a least squares problem, we first remove the constraint on the test matrix $\mathbf{X}$, replacing $\mathbf{X}$ with $\mathbf{P}_{\mathbf{U}}^{\perp}\mathbf{X}$ where $\mathbf{P}_{\mathbf{U}}^{\perp}$ is the orthogonal projector onto the complement of the range of $\mathbf{U}$, $\mathbf{P}_{\mathbf{U}}^{\perp} = \mathbf{I} - \mathbf{U}\mathbf{U}^{\top}$. This leaves a linear system of equations over all test matrices $\mathbf{X} \in \mathbb{R}^{m\times n}$

$$(\operatorname{vec}\mathbf{P}_{\mathbf{U}}^{\perp}\mathbf{X})^{\top}(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))^{\top}(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))(\operatorname{vec}\boldsymbol{\Delta}) = (\operatorname{vec}\mathbf{P}_{\mathbf{U}}^{\perp}\mathbf{X})^{\top}(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))^{\top}\mathbf{r}(\mathbf{U}).$$

As before, the projector $\mathbf{P}_{\mathbf{U}}^{\perp}$ vanishes by Theorem 1

$$\begin{aligned}\left[(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))(\operatorname{vec}\mathbf{P}_{\mathbf{U}}^{\perp}\mathbf{X})\right]_i &= \mathsf{Tr}\left[(\mathbf{P}_{\mathbf{U}}^{\perp}\mathbf{X})^{\top}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}\right] = \mathsf{Tr}\left[\mathbf{X}^{\top}(\mathbf{I}-\mathbf{U}\mathbf{U}^{\top})[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}\right] \\ &= \mathsf{Tr}\left[\mathbf{X}^{\top}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{i,\cdot,\cdot}\right] = (\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))(\operatorname{vec}\mathbf{X}). \end{aligned} \tag{30}$$

If we use the coordinate matrices $\mathbf{e}_i\mathbf{e}_j^{\top}$ as a basis for $\mathbf{X} \in \mathbb{R}^{m\times n}$, we then recover the normal equations

$$(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))^{\top}(\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))(\operatorname{vec}\boldsymbol{\Delta}) = (\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U}))^{\top}\mathbf{r}(\mathbf{U}). \tag{31}$$

Hence as in a Euclidean space (c.f. [27, (10.26)]), the Gauss-Newton step $\boldsymbol{\Delta}$ is the solution to the linear least squares problem

$$\underset{\substack{\boldsymbol{\Delta}\in\mathbb{R}^{m\times n}\\ \mathbf{U}^{\top}\boldsymbol{\Delta}=\mathbf{0}}}{\text{minimize}} \left\|\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})\operatorname{vec}\boldsymbol{\Delta} + \mathbf{r}(\mathbf{U})\right\|_2^2. \tag{32}$$

Finally, we note that this linear least squares problem (32) can be solved using the pseudoinverse and automatically satisfy the constraint $\mathbf{U}^{\top}\boldsymbol{\Delta} = \mathbf{0}$. Using (30) we can insert the projector $\mathbf{P}_{\mathbf{U}}^{\perp}$ into the Gauss-Newton step (32)

$$\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})\operatorname{vec}\boldsymbol{\Delta} = \operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})\operatorname{vec}(\mathbf{P}_{\mathbf{U}}^{\perp}\boldsymbol{\Delta}) = \operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})[\mathbf{I}_n \otimes \mathbf{P}_{\mathbf{U}}^{\perp}]\operatorname{vec}\boldsymbol{\Delta} \tag{33}$$

where $\otimes$ is the Kronecker product. Hence $\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})$ has a nullspace of dimension n^2 which contains $\mathbf{I}_n \otimes \mathbf{U}\mathbf{U}^{\top}$. Thus, if (32) is solved via the pseudoinverse, $\boldsymbol{\Delta}$ will automatically obey the constraint $\mathbf{U}^{\top}\boldsymbol{\Delta} = \mathbf{0}$:

$$\operatorname{vec}\Delta = -[\operatorname{vec}\boldsymbol{\mathcal{J}}(\mathbf{U})]^{+}\mathbf{r}(\mathbf{U}) = -\begin{bmatrix}[\boldsymbol{\mathcal{J}}(\mathbf{U})]_{\cdot,\cdot,1} & \cdots & [\boldsymbol{\mathcal{J}}(\mathbf{U})]_{\cdot,\cdot,n}\end{bmatrix}^{+}\mathbf{r}(\mathbf{U}). \tag{34}$$

The asymptotic cost of constructing $\boldsymbol{\Delta}$ via the pseudoinverse is the same as via the normal equations: an $\mathcal{O}(N(mn)^2)$ operation SVD where $N > mn$ as opposed to an $\mathcal{O}((mn)^3)$ dense linear solve. However the pseudoinverse solution is better conditioned, not suffering from the squaring of the condition number as in the normal equations [4, §2.3.3].

3.3. Algorithm. We now combine the Gauss-Newton step (34) with backtracking along the geodesic (14) to construct a convergent data-driven polynomial ridge approximation algorithm. The complete algorithm is given in Algorithm 1, using a pseudo-inverse to construct the Gauss-Newton step as in (34). We ensure convergence by inserting a check on line 12 to ensure $\boldsymbol{\Delta}$ is always a decent direction, and if not, replacing it with the negative gradient. Then the sequence of $\boldsymbol{\Delta}$ is a *gradient related sequence* [1, Def. 4.2.1] and the iterates $\mathbf{U}$ converge to a stationary point where $\operatorname{grad}\phi(\mathbf{U}) = \mathbf{0}$ by [1, Cor. 4.3.2] since the Grassmann manifold is compact [37, §9].

In practice, we have found choosing a small Armijo tolerance β, such as $\beta = 10^{-6}$, helps speed convergence. Our implementation equips this algorithm with three different convergence criteria: small change between $\mathbf{U}$ and $\mathbf{U}_+$ as measured by the smallest canonical subspace angle, small change in the norm of the residual, and small gradient norm.

Algorithm 1 Variable projection polynomial ridge approximation

Input : Sample points $\mathbf{X} \in \mathbb{R}^{N\times m}$; function values $\mathbf{f} \in \mathbb{R}^N$, $[\mathbf{f}]_i = f([\mathbf{X}]_{i,\cdot})$; subspace dimension n; polynomial degree p (if $p = 1$, then $n = 1$); step length reduction factor $\gamma \in (0, 1)$; Armijo tolerance $\beta \in (0, 1)$.

Output : Active subspace $\mathbf{U} \in \mathbb{R}^{m\times n}$; polynomial coefficients $\mathbf{c} \in \mathbb{R}^M$, $M = \binom{n+p}{p}$.

Sample entries of $\mathbf{Z} \in \mathbb{R}^{m\times n}$ from a normal distribution;
Compute short form QR $\mathbf{UR} \leftarrow \mathbf{Z}$ *hence* $\mathbf{U}$ *uniformly samples* $\mathbb{G}(n, \mathbb{R}^m)$
repeat
  Build $\mathbf{V}(\mathbf{U})$ following (5): $\mathbf{V} \leftarrow \mathbf{V}(\mathbf{U})$;
  Compute polynomial coefficients $\mathbf{c} \leftarrow \mathbf{V}^+\mathbf{f}$;
  Compute the residual: $\mathbf{r} \leftarrow \mathbf{r}(\mathbf{U}) = \mathbf{f} - \mathbf{Vc}$;
  Build the Jacobian following (10): $\boldsymbol{\mathcal{J}} \leftarrow \boldsymbol{\mathcal{J}}(\mathbf{U}) \in \mathbb{R}^{N\times m\times n}$;
  Build the gradient following (29): $\mathbf{G} \leftarrow \mathbf{G}(\mathbf{U}) = \sum_{i=1}^N [\boldsymbol{\mathcal{J}}]_{i,\cdot,\cdot}[\mathbf{r}]_i$;
  Compute the short form SVD: $\mathbf{Y\Sigma Z}^\top \leftarrow \operatorname{vec} \boldsymbol{\mathcal{J}}$;
  Compute the Gauss-Newton step following (34): $\operatorname{vec}\boldsymbol{\Delta} \leftarrow -\sum_{i=1}^{M-n^2} [\boldsymbol{\Sigma}]_{i,i}^{-1}[\mathbf{Z}]_{\cdot,i}[\mathbf{Y}^\top\mathbf{r}]_i$;
  Compute slope along Gauss-Newton step: $\alpha \leftarrow \operatorname{Tr} \mathbf{G}^\top\boldsymbol{\Delta} = (\operatorname{vec}\mathbf{G})^\top(\operatorname{vec}\boldsymbol{\Delta})$;
  if $\alpha \geq 0$ **then** $\boldsymbol{\Delta}$ *is not a descent direction*
    $\boldsymbol{\Delta} \leftarrow \mathbf{G}$;
    $\alpha \leftarrow \operatorname{Tr} \mathbf{G}^\top\boldsymbol{\Delta}$;
  Compute the short form SVD: $\mathbf{Y\Sigma Z}^\top \leftarrow \boldsymbol{\Delta}$;
  for $t = \gamma^0, \gamma^1, \gamma^2, \ldots$ **do** *backtracking line search*
    Compute new step following (14): $\mathbf{U}_+ \leftarrow \mathbf{UZ}\cos(\boldsymbol{\Sigma}t)\mathbf{Z}^\top + \mathbf{Y}\sin(\boldsymbol{\Sigma}t)\mathbf{Z}^\top$;
    Compute new residual: $\mathbf{r}_+ \leftarrow \mathbf{f} - \mathbf{V}(\mathbf{U}_+)\mathbf{V}(\mathbf{U}_+)^+\mathbf{f}$;
    if $\|\mathbf{r}_+\|_2 \leq \|\mathbf{r}\|_2 + \alpha\beta t$ **then** *Armijo condition satisfied*
      break;
  Update the estimate $\mathbf{U} \leftarrow \mathbf{U}_+$;
until $\mathbf{U}$ *converges*;

4. Examples. To demonstrate the effectiveness of the proposed algorithm for ridge approximation, Algorithm 1, we apply it to a mixture of synthetic and application problems. Code generating these examples along with an implementation of Algorithm 1 are available at `https://github.com/jeffrey-hokanson/varproridge`. First we compare the proposed algorithm to the alternating approach of [9] on a problem where the ridge function is known a priori and examine both convergence and wall clock time. Next, we demonstrate that even though finding the ridge approximation is a non-convex problem, the proposed algorithm frequently finds the global minimizer from a random initialization. Then we illustrate how the discrete measure induced by the point set $\{\mathbf{x}_i\}_{i=1}^N$ can change the resulting ridge approximation subspace $\mathbf{U}$. Finally, we apply the proposed algorithm to two application problems: modeling lift and drag from a NACA0012 airfoil with 18 parameters and modeling mean solution on the Neumann boundary of an elliptic PDE with 100-random coefficients. In both cases we are able to find high-quality ridge approximations using the proposed algorithm.

4.1. Convergence. As a first test, we compare the performance of the proposed Gauss-Newton based algorithm to the alternating algorithm of [9] for problems with and without noise. Following the analysis of Ruhe and Wedin [32, §3] we expect the proposed variable projection Gauss-Newton approach (in their notation, Algorithm

I) to converge at the rate of Gauss-Newton and the alternating approach (in their notation, Algorithm III) to converge linearly. Hence, for problems with no residual, we expect the proposed approach to converge quadratically and linearly otherwise [27, §10.2].

Our numerical experiments bare out this analysis. In Figure 2 we compare the performance of the alternating and Gauss-Newton approaches. In this example, f is taken to be a cubic ridge function on a two dimensional subspace:

$$f : \mathcal{D} \subset \mathbb{R}^{10} \to \mathbb{R}; \quad f(\mathbf{x}) = (\mathbf{e}_1^\top \mathbf{x})^2 + (\mathbf{1}^\top \mathbf{x}/10)^3 + 1; \quad \mathcal{D} = [-1, 1]^{10}. \tag{35}$$

Sampling $N = 1000$ points uniformly over the domain, Figure 2 shows the per-iteration convergence history for ten different initializations of each algorithm. As expected in the case with no residual, the proposed Gauss-Newton approach converges quadratically while the alternating approach only converges linearly. In the bottom row of this figure, Gaussian random noise with unit variance has been added to the function values $f(\mathbf{x}_i)$. Although in this case we can no longer guarantee quadratic convergence, the proposed approach rapidly converges where as the alternating approach requires more iterations on average.

This figure also exposes a numerical issue in the construction of the alternating algorithm of [9]. Their algorithm alternates between finding polynomial coefficients given $\mathbf{U}$ via linear least squares and optimizing $\mathbf{U}$ with the coefficients $\mathbf{c}$ fixed using `pymanopt` [35] applied to the squared residual; i.e.,

$$\underset{\text{Range } \mathbf{U} \in \mathbb{G}(m, \mathbb{R}^n)}{\text{minimize}} \ \frac{1}{2} \|\mathbf{f} - \mathbf{V}(\mathbf{U})\mathbf{c}\|_2^2. \tag{36}$$

In their implementation, `pymanopt` solves this optimization problem using gradient descent. We suspect that the early plateau in the residual is due to numerical issues as the plateau does not decrease if the termination criteria are made more strict.

4.2. Performance. To asses the overall performance of both the alternating implementation from [9] and the proposed Gauss-Newton approach, we compared the wall clock performance of both algorithms fitting $N = 1000$ samples of

$$f_{n,p} : \mathcal{D} \subset \mathbb{R}^{10} \to \mathbb{R}; \quad f_{n,p}(\mathbf{x}) = (\mathbf{1}^\top \mathbf{x})^p + \sum_{j=1}^{m-1} (\mathbf{e}_j^\top \mathbf{x})^{p-1}; \quad \mathcal{D} = [-1, 1]^{10}. \tag{37}$$

Table 1 shows the wall clock time taken by both algorithms for a range of polynomial degrees p and subspace dimensions m. In each case, both algorithms were asked to fit a ridge approximation to identical data from identical starting conditions. To ensure a fair comparison, both algorithms were stopped when the normalized residual dropped below 10^{-5}, which both algorithms were capable of in Figure 2. Overall, the wall clock time for our proposed algorithm is about $4000\times$ faster than the alternating approach. Part of this improvement is due to the algorithms themselves. Gauss-Newton's quadratic convergence requires fewer iterations than a linearly convergent alternating method and each iteration of the alternating method is more expensive as it consists of consists of multiple gradient descent steps. However, another significant part of the improved performance is due to more efficient code that exploits the specific structure of this problem, removing several large matrix-matrix products.

4.3. Convergence to global minimizer. One concern is that our proposed algorithm might converge to a spurious local minimizer rather than the global minimizer since ridge approximation (2) is not necessarily a convex problem. However,

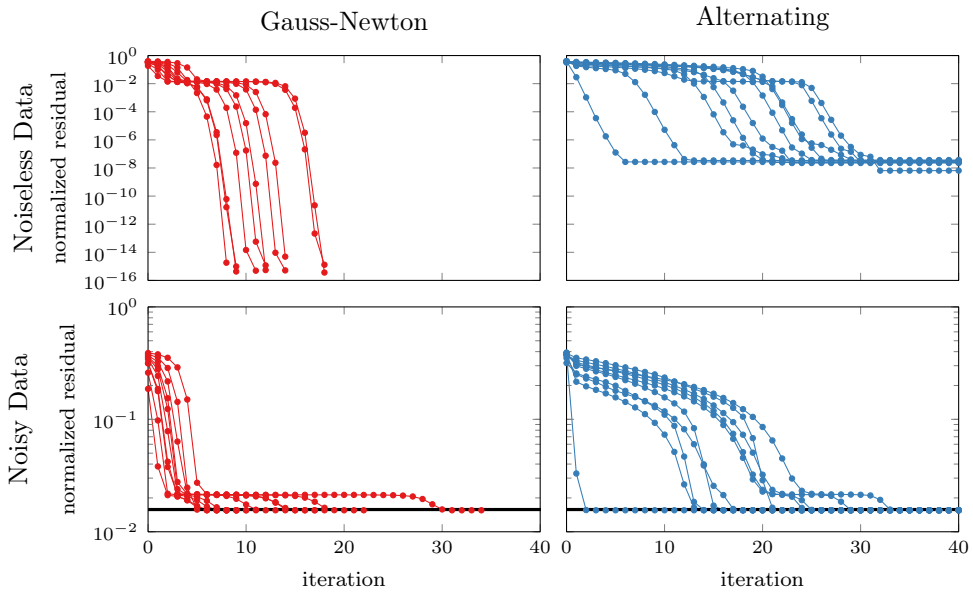

FIG. 2. *A comparison of the per-iteration performance of our Gauss-Newton method (Algorithm 1) and the alternating approach of Constantine, Eftekhari, and Ward [9, Alg. 2]. The top rows compare these algorithms when finding a two variable cubic ridge function approximation from* $N = 1000$ *uniform random samples of* (35). *The bottom row shows the performance when Gaussian random noise was added to each sample of* f *to simulate a function that is not an exact ridge function; the black line indicates the norm of the noise introduced.*

TABLE 1

Median wall clock time in seconds from ten replicates of each algorithm applied to identical zero-residual data consisting of $N = 1000$ *uniform samples of* $f_{n,p}$ *from (37), initialized with the same random subspace* $\mathbf{U}$*, and stopped when the normalized residual reached* 10^{-5}*. Experiments were conducted on a 2013 Mac Pro with a six-core Intel Xeon CPU E5-1650 v2 clocked at 3.50GHz with 16GB RAM.*

p	Gauss-Newton					Alternating				
	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$
2	0.009	0.013	0.026	0.035	0.044	66.5	90.5	25.6	32.9	42.5
3	0.010	0.025	0.033	0.067	0.108	51.9	142.3	181.7	295.8	381.8
4	0.010	0.025	0.063	0.176	0.364	69.9	237.6	473.9	1,065.8	1,780.5
5	0.012	0.034	0.118	0.431	1.082	98.2	364.7	962.2	2,754.9	4,348.4

our numerical results demonstrate that this algorithm frequently converges to the global minimizer regardless of initial subspace estimate. To provide an example of this, we construct a quadratic ridge approximation on an n-dimensional subspace from $N = 1000$ samples of

$$f_n : \mathcal{D} \subset \mathbb{R}^{10} \to \mathbb{R}; \quad f_n(\mathbf{x}) = \sum_{j=1}^{n} (\mathbf{e}_j^\top \mathbf{x})^2; \quad \mathcal{D} = [-1, 1]^{10}. \tag{38}$$

As we know the true subspace and that the residual should be near machine precision, we can easily assess if the algorithm has not correctly converged. Table 2 shows the frequency with which our proposed algorithm terminated with an incorrect subspace. In general, this frequency is low and so it should be sufficient in many practical cases

TABLE 2
Probability of not finding the global minimizer given a random initial subspace of dimension n. This probability was estimated from 1000 trials of Algorithm 1 fitting a quadratic ridge function of n variables to $N = 1000$ samples of $f(\mathbf{x})$ from (38) taken uniformly over the domain.

$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$	$n=7$	$n=8$	$n=9$	$n=10$
0.0%	10.8%	16.2%	15.2%	7.3%	7.6%	10.6%	7.1%	3.4%	0.0%

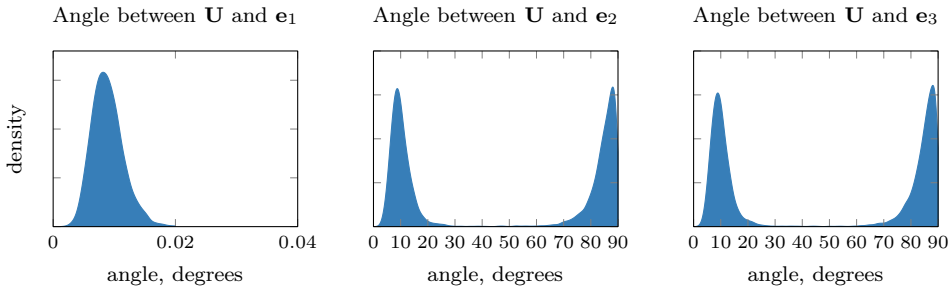

FIG. 3. *Angles between the subspace recovered via Algorithm 1 from different point sets $\{\mathbf{x}_i\}_{i=1}^{1000}$ uniformly sampling $\mathcal{D}$ from (39). Each plot shows a kernel density estimate of the probability density function for the subspace angle between $\mathbf{U}$ and the unit vector $\mathbf{e}_k$ using 10^4 samples.*

to try multiple random initializations and take the best to ensure convergence to an approximate global minimizer.

4.4. Consistency. An important factor to consider when approximating functions from discrete point set is the approximation may depend greatly on which points $\{\mathbf{x}_i\}_{i=1}^N$ are used to construct the approximation. Consider, for example, the following ridge quadratic function on a three-dimensional subspace

$$f : \mathcal{D} \subset \mathbb{R}^{10} \to \mathbb{R}; \quad f(\mathbf{x}) = 1000[\mathbf{x}]_1^2 + [\mathbf{x}]_2^2 + [\mathbf{x}]_3^2; \quad \mathcal{D} = [-1,1]^{10}. \tag{39}$$

If we seek to approximate this function with a quadratic ridge function on a *two-dimensional subspace*, we expect that $\mathbf{U}$ would (i) contain $\mathbf{e}_1$ due to the large weighting and (ii) choose between either $\mathbf{e}_2$ and $\mathbf{e}_3$. The choice between these two may depend on which $\mathbf{x}_i$ were sampled; i.e., if there were more points along on the boundary in the $\mathbf{e}_2$ direction than $\mathbf{e}_3$, then the subspace would include $\mathbf{e}_2$ rather than $\mathbf{e}_3$. Figure 3 illustrates this behavior for the f from (39). As expected, the recovered subspace $\mathbf{U}$ contains $\mathbf{e}_1$ to high precision, but alternates between $\mathbf{e}_2$ and $\mathbf{e}_3$ depending on the samples $\{\mathbf{x}_i\}_{i=1}^N$.

4.5. NACA airfoil. As a first demonstration of our proposed algorithm on an application problem, we consider an 18-parameter model of a NACA0012 airfoil with two quantities of interest: lift and drag. This model from [6, §5.3.1] consists 18 Hicks-Henne parameters that modify the airfoil geometry and the quantity of interest is computed using the Stanford University Unstructured (SU2) computational fluid dynamics code [14]. Table 3 reports on the ability of a polynomial ridge function model to fit this problem. As these results show, for this problem a polynomial degree greater than 2 does not substantially improve the fit; however, increasing the subspace dimension can yield gradual improvement.

TABLE 3
Normalized residual $\|\mathbf{f}-\mathbf{g}\|_2/\|\mathbf{f}\|_2$ *where* $[\mathbf{g}]_i = g(\mathbf{U}^\top \mathbf{x}_i)$ *for* $N = 1756$ *samples of the NACA0012 airfoil for a variety of polynomial degrees* p *and subspace dimensions* m.

Lift
constant: 0.48802, linear: 0.12051

dim	$p=2$	$p=3$	$p=4$	$p=5$
1	0.12026	0.12016	0.11958	0.11925
2	0.08870	0.08609	0.08459	0.08359
3	0.07515	0.07106	0.06846	0.06587
4	0.06894	0.06175	0.05767	0.05351
5	0.06668	0.05396	0.04826	0.04177
6	0.06470	0.04897	0.04162	0.03336
7	0.06305	0.04529	0.03504	0.02428
8	0.06217	0.04189	0.02933	0.01857
9	0.06139	0.03899	0.02315	

Drag
constant: 0.44815, linear: 0.17496

dim	$p=2$	$p=3$	$p=4$	$p=5$
1	0.16100	0.15955	0.15949	0.15938
2	0.11818	0.11409	0.11281	0.11189
3	0.09464	0.08612	0.08404	0.08221
4	0.07378	0.05971	0.05662	0.05258
5	0.06977	0.04969	0.04388	0.03859
6	0.06919	0.04203	0.03370	0.02603
7	0.06781	0.03891	0.02841	0.01769
8	0.06741	0.03627	0.02324	0.01278
9	0.06706	0.03419	0.01842	

TABLE 4
Normalized residual $\|\mathbf{f}-\mathbf{g}\|_2/\|\mathbf{f}\|_2$ *where* $[\mathbf{g}]_i = g(\mathbf{U}^\top \mathbf{x}_i)$ *for* $N = 10^4$ *samples of the elliptic PDE* (40) *for a variety of polynomial degrees* p *and subspace dimensions* m.

Short
constant: 0.06301, linear: 0.00971

dim	$p=2$	$p=3$	$p=4$	$p=5$
1	0.00938	0.00938	0.00938	0.00938
2	0.00880	0.00879	0.00879	0.00880
3	0.00812	0.00812	0.00811	0.00823
4	0.00759	0.00758	0.00752	0.00791
5	0.00709	0.00696	0.00693	0.00688
6	0.00659	0.00656	0.00642	0.00640
7	0.00609	0.00604	0.00596	0.00603
8	0.00575	0.00559	0.00547	0.00673
9	0.00541	0.00530	0.00513	0.00636
10	0.00505	0.00491	0.00465	0.00541

Long
constant: 0.72424, linear: 0.41477

dim	$p=2$	$p=3$	$p=4$	$p=5$
1	0.19757	0.08209	0.03964	0.02947
2	0.19200	0.06806	0.02995	0.02313
3	0.18843	0.06095	0.02580	0.01834
4	0.18634	0.05570	0.02122	0.01290
5	0.18358	0.05217	0.01709	0.00927
6	0.18162	0.04947	0.01463	0.00719
7	0.17969	0.04711	0.01248	0.00577
8	0.17823	0.04491	0.01117	0.00484
9	0.17685	0.04259	0.00998	0.00425
10	0.17524	0.04084	0.00868	0.00374

4.6. PDE. As a final demonstration of the algorithm, we consider a 2D elliptic PDE from [8] with 100 input parameters characterizing the operator coefficients:

$$-\nabla_{\mathbf{s}} \cdot (a \nabla_{\mathbf{s}} u) = 1, \quad \mathbf{s} \in [0,1]^2, \tag{40}$$

where $a = a(\mathbf{s}, \mathbf{x})$ with $\mathbf{x} \in \mathbb{R}^{100}$. The parameters $\mathbf{x}$ are coefficients in a 100-term truncated Karhunen-Loéve expansion of $\log(a)$ with a correlation function,

$$\mathrm{Corr}(\mathbf{s}_1, \mathbf{s}_2) \;=\; \exp\left(\frac{-\|\mathbf{s}_1 - \mathbf{s}_2\|_1}{2\ell}\right), \tag{41}$$

where ℓ is a correlation length parameter. The boundary conditions are homogeneous Neumann conditions on the right boundary and homogeneous Dirichlet conditions on the other boundaries. The output of interest is the spatial average of the solution u on the right boundary. Table 4 shows the normalized residual of the ridge approximation for two cases of interest: (i) a long correlation length $\ell = 1$ and (ii) a short correlation length $\ell = 0.01$.

5. Summary and discussion. We have proposed and studied a Gauss-Newton optimization method on the Grassmann manifold for a variable projection reformulation of the data-driven polynomial ridge approximation problem. The low-dimensional

structure in the ridge approximation enables dimension reduction that leads to increased efficiency in parameter studies, e.g., optimization and uncertainty quantification. Our Gauss-Newton approach for building ridge approximations has theoretical advantages (e.g., faster convergence rates) and demonstrates superior numerical performance compared to the alternating approach that switches between minimizing over polynomial coefficients and subspaces.

Our focus has been to solve the data-drive polynomial ridge approximation problem efficiently and accurately, however, several open questions remain. For example, how should we select the polynomial degree p and subspace dimension n for the ridge approximation? If $f(\mathbf{x})$ contains random noise with a known distribution as in statistical regression, then approaches to choose p and n include Akaike information criterion (AIC) [2] or cross-validation [30]. However, it is unclear how to interpret either of these approaches when $f(\mathbf{x})$ is deterministic or is corrupted by deterministic computational errors.

Another open question is how to iteratively choose new samples $\mathbf{x}_i$ to optimally fit the ridge approximation model? This applies regardless of optimization method. Here, we have assumed these samples were given; however, random samples in the high-dimensional design space are unlikely to capture the full range of the function. One idea is to first construct g and $\mathbf{U}$ from the initial samples and then to use the recovered g and $\mathbf{U}$ to choose subsequent samples of $\mathbf{x}$, repeating as necessary. This approach was taken by Cohen et al. [5] in the case of a one-dimensional ridge function where $\mathbf{U}$ had all positive entries. An open question is how these results may extend to the multivariate polynomial ridge approximation case.

A final question is how to improve the numerical stability of this algorithm. Vandermonde matrices are well known to be ill-conditioned [21, Ch. 22] and hence our algorithm is likely to accumulate errors when the polynomial degree is large. In one dimension this can be alleviated using the high relative accuracy SVD algorithm of Demmel [12, §5], however an open question is if these results can be extended to the multidimensional Vandermonde matrix used by our algorithm.

REFERENCES

[1] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, 2008.

[2] H. Akaike, *A new look at the statistical model identification*, IEEE T. Automat. Contr., 19 (1974), pp. 716–723, https://doi.org/10.1109/TAC.1974.1100705.

[3] M. Bauerheim, A. Ndiaye, P. Constantine, S. Moreau, and F. Nicoud, *Symmetry breaking of azimuthal thermoacoustic modes: the UQ perspective*, J. Fluid Mech., 789 (2016), pp. 534–566, https://doi.org/10.1017/jfm.2015.730.

[4] Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, PA, 1996.

[5] A. Cohen, I. Daubechies, R. DeVore, G. Kerkyacharian, and D. Picard, *Capturing ridge functions in high dimensions from point queries*, Constr. Approx., 35 (2012), pp. 225–243, https://doi.org/10.1007/s00365-011-9147-6.

[6] P. G. Constantine, *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*, SIAM, Philadelphia, 2015.

[7] P. G. Constantine and A. Doostan, *Time-dependent global sensitivity analysis with active subspaces for a lithium ion battery model*, 2016, https://arxiv.org/abs/1606.08770.

[8] P. G. Constantine, E. Dow, and Q. Wang, *Active subspace methods in theory and practice: Applications to kriging surfaces*, SIAM J. Sci. Comp., 36 (2014), pp. A1500–A1524, https://doi.org/10.1137/130916138.

[9] P. G. Constantine, A. Eftekhari, and R. A. Ward, *A near-stationary subspace for ridge approximation*, 2016, https://arxiv.org/abs/arXiv:1606.01929.

[10] P. G. Constantine, M. Emory, J. Larsson, and G. Iaccarino, *Exploiting active subspaces to quantify uncertainty in the numerical simulation of the HyShot II scramjet*, J. Comput.

Phys., 302 (2015), pp. 1–20, https://doi.org/10.1016/j.jcp.2015.09.001.
[11] P. G. Constantine, B. Zaharatos, and M. Campanelli, *Discovering an active subspace in a single-diode solar cell model*, Statistical Analysis and Data Mining: The ASA Data Science Journal, 8 (2015), pp. 264–273, https://doi.org/10.1002/sam.11281.
[12] J. Demmel, *Accurate singular value decompositions of structured matrices*, SIAM J. Matrix Analy. A., 21 (1999), pp. 562–580.
[13] P. Diaconis and M. Shahshahani, *On nonlinear functions of linear combinations*, SIAM J. Sci. Stat. Comp., 5 (1984), pp. 175–191, https://doi.org/10.1137/0905013.
[14] T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso, *SU2: An open-source suite for multiphysics simulation and design*, AIAA J., 54 (2016), pp. 828–846, https://doi.org/10.2514/1.J053813.
[15] A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Analy. A., 20 (1998), pp. 303 – 353, https://doi.org/10.1137/S0895479895290954.
[16] M. Fornasier, K. Schnass, and J. Vybiral, *Learning functions of few arbitrary linear parameters in high dimensions*, Found. Comput. Math., 12 (2012), pp. 229–262, https://doi.org/10.1007/s10208-012-9115-y.
[17] J. H. Friedman and W. Stuetzle, *Projection pursuit regression*, J. Am. Stat. Assoc., 76 (1981), pp. 817–823, https://doi.org/10.1080/01621459.1981.10477729.
[18] J. M. Gilbert, J. L. Jefferson, P. G. Constantine, and R. M. Maxwell, *Global spatial sensitivity of runoff to subsurface permeability using the active subspace method*, Adv. Water Resour., 92 (2016), pp. 30–42, https://doi.org/10.1016/j.advwatres.2016.03.020.
[19] A. Glaws, P. G. Constantine, J. Shadid, and T. M. Wildey, *Dimension reduction in MHD power generation models: dimensional analysis and active subspaces*, 2016, https://arxiv.org/abs/1609.01255.
[20] G. H. Golub and V. Pereyra, *The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate*, SIAM Journal on Numerical Analysis, 10 (1973), pp. 413–432.
[21] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, PA, 2002.
[22] J. L. Jefferson, J. M. Gilbert, P. G. Constantine, and R. M. Maxwell, *Active subspaces for sensitivity analysis and dimension reduction of an integrated hydrologic model*, Comput. Geosci., 83 (2015), pp. 127–138, https://doi.org/10.1016/j.cageo.2015.11.002.
[23] D. R. Jones, *A taxonomy of global optimization methods based on response surfaces*, J. Global Optim., 21 (2001), pp. 345–385, https://doi.org/10.1023/A:1012771025575.
[24] O. C. Lingjærde and K. Liestøl, *Generalized projection pursuit regression*, SIAM J. Sci. Comp., 20 (1998), pp. 844–857, https://doi.org/10.1137/S1064827595296574.
[25] T. W. Lukaczyk, *Surrogate modeling and active subspaces for efficient optimization of supersonic aircraft*, PhD thesis, Stanford University, 2015, https://purl.stanford.edu/xx611nd3190.
[26] T. W. Lukaczyk, P. Constantine, F. Palacios, and J. J. Alonso, *Active subspaces for shape optimization*, 10th AIAA Multidisciplinary Design Optimization Conference, (2014), https://doi.org/10.2514/6.2014-1171.
[27] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer, New York, 2006.
[28] E. Novak and H. Woźniakowski, *Approximation of infinitely differentiable multivariate functions is intractable*, J. Complexity, 25 (2009), pp. 398–404, https://doi.org/10.1016/j.jco.2008.11.002.
[29] C. Othmer, T. W. Lukaczyk, P. Constantine, and J. J. Alonso, *On active subspaces in car aerodynamics*, 17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, (2016), https://doi.org/10.2514/6.2016-4294.
[30] R. R. Picard and R. D. Cook, *Cross-validation of regression models*, J. Am. Stat. Assoc., 79 (1984), pp. 575–583, https://doi.org/10.1080/01621459.1984.10478083.
[31] C. B. Roosen and T. J. Hastie, *Automatic smoothing spline projection pursuit*, J. Comput. Graph. Stat., 3 (1994), pp. 235–248, https://doi.org/10.1080/10618600.1994.10474642.
[32] A. Ruhe and P. A. Wedin, *Algorithms for separable nonlinear least squares problems*, SIAM Rev., 22 (1980), pp. 318–337, https://doi.org/10.1137/1022057.
[33] S. Shan and G. G. Wang, *Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions*, Struct. Multidisc. Optim., 41 (2010), pp. 291–241, https://doi.org/10.1007/s00158-009-0420-2.
[34] R. C. Smith, *Uncertainty Quantificaation: Theory, Implementation, and Applications*, SIAM, Philidelphia, PA, 2014.

[35] J. Townsend, N. Koep, and S. Weichwald, *Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation*, J. Mach. Learn. Res., 17 (2016), pp. 1–5, http://jmlr.org/papers/v17/16-177.html.

[36] H. Tyagi and V. Cevher, *Learning non-parametric basis independent models from point queries via low-rank methods*, Appl. Comput. Harmon. Anal., 37 (2014), pp. 389–412, https://doi.org/http://dx.doi.org/10.1016/j.acha.2014.01.002.

[37] Y.-C. Wong, *Differential geometry of Grassmann manifolds*, Proc. Nat. Acad. Sci. U.S.A., 57 (1967), pp. 589–594, https://doi.org/10.1073/pnas.57.3.589.